

REMARKS

At present, Applicants' claims 1-4 and 9-20 stand rejected under 35 USC § 102(e) based upon the published patent application of Moore et al. (U.S. Patent Publication US 2004/0249904 A1, having a publication date of December 9, 2004 and bearing a filing date of April 16, 2003, hereinafter referred to simply as Moore). Claims 5-8 also stand rejected under 35 USC § 103 based upon the aforementioned Moore application in further review of the published patent application of Dugan et al. (U.S. Patent Publication US 2006/0165223 A1, having a publication date of July 27, 2006 and bearing a filing date of February 15, 2006, hereinafter referred to simply as Dugan). Applicants respectfully, but most strenuously, traverse all of the rejections for the reasons set forth below.

35 USC § 102 Rejection

With respect to the rejection under 35 USC § 102, the Examiner is again reminded of the fact that rejections under this statute constitute a narrow ground of rejection. The so-called anticipation rejection requires each and every recited claim element to be found within the four corners of a single document. Any of the minor exceptions to this rule are not germane to the present discussion. Attention is now focused on a discussion of the differences found between the patent application of Moore and the subject claim language.

The Examiner's attention is specifically directed to the following language found in Applicants' claim 1 and the claims which depend there from. Is also to be particularly noted that language to the same effect is also to be found in claim 19 and 20.

“posting a worker thread to one or more of the nodes to perform data movement in response to the event.”

There is no teaching disclosure or suggestion in either of the two documents cited against applicants claims relating to the posting of a worker thread. In particular, it is noted that the published patent application of Moore, while it references “threads” in several places, nowhere does this document teach the posting of a thread as part of their process. The following quotation from the Moore publication is particularly telling in this respect

“[0117] Preferably interruptible token acquisition is used to enable recovery and relocation in several ways: (1) threads processing messages from failed nodes that are waiting for the token state to stabilize are sent an interrupt to be terminated to allow recovery to begin; (2) threads processing messages from failed nodes which may have initiated a token recall and are waiting for the tokens to come back are interrupted; (3) threads that are attempting to lend tokens which are waiting for the token state to stabilize and are blocking recovery/relocation are interrupted; and (4) threads that are waiting for the token state to stabilize in a filesystem that has been forced offline due to error are interrupted early. Threads waiting for the token state to stabilize first call a function to determine if they are allowed to wait, i.e. none of the factors above apply, then go to sleep until some other thread signals a change in token state.”

In this portion of their document Moore et al. discuss the use of threads but never do they teach the posting of a worker thread. Below is summarized their teachings with respect to threads as specified in paragraph [0117]:

threads...are sent an interrupt....

threads...are interrupted...

threads...are interrupted...

threads...are interrupted early.

However it is noted that nowhere is there mentioned the posting of a thread. Furthermore, there is nothing recited in the publication of Moore et al. which describes anything which could be considered as a “worker thread,” as specifically recited in all Applicants’ claims.

In another portion of the cited document, there is also a reference to “threads.” This is seen in the quotations below from the cited document:

“[0103] An RPC is a thread initiated on a node in response to a message from another node to act as a proxy for the requesting node.”

“[0105] When celldown callouts have been performed 268 for all of the objects associated with a failed node, the operations frozen 266 previously are thawed or released 270. The message channel is thawed 270, so that any threads that are waiting for responses can receive error messages that a cell is down, i.e., a node has

failed, so that that the threads can do any necessary cleanup and then drop the handle hold.” [See the Moore et al. document for their use of the word “chandle.”]

This material, along with the material found in paragraphs [0103] and [0104] is also useful to consider since the Examiner has cited it several times and seems to be significantly relying upon it. These lines are part of the recovery and relocation section of the Moore et al. application starting at paragraph [0089] and illustrated in Figures 10-12. The situation is that a member of their cluster, possibly including a metadata server has failed. In this case, one of their nodes is elected a leader for the recovery (paragraph [0092]) and must perform certain specific recovery tasks including those described in [0103]-[0105]. Specifically, [0103]-[0105] involve calling each surviving member of the cluster and repairing their lock state to reflect the state of the cluster after the failure. **This task has nothing whatsoever to do with data movement within the context of a DMAPI event.**

Applicant’s attorney has also undertaken an exhaustive effort to identify all of the locations within the cited document which refers to the notion of a thread. Accordingly, one also finds the quotation below from Moore et al.

“[0118] To interrupt, CORPSE and KORE each wake all sleeping threads. These threads loop, check if the token state has changed and if not attempt to go back to sleep. This time, one of the factors above may apply and if so a thread discovering it returns immediately with an “early” status...In the token recall case, this means the thread will have to leave the token server data structure in a partially recalled state. This transitory state is exited when the last of the recalls comes in, and the thread returning the last recalled token clears the state. In lending cases, the thread will return early, potentially without all tokens desired for lending.”

Again it is pointed out to the examiner that all of the rejected claims specifically include the recitation of “posting a worker thread to one or more of the nodes to perform data movement in response to the event.” Note that this is within the context of the DMAPI event otherwise also recited in applicants claim 1. The whole purpose of the present invention is to improve the performance of event handling for DMAPI events. The cited portions [0103]-[0105] above deal with locking in the recovery of the failure of a machine as has been clearly indicated above. Portions [0117]-[0118] provide more detail on lock recovery. **However, there is nothing**

contained within this material which would in any way teach disclose or suggest a process of data movement within the context of a DMAPI event.

The Examiner has also asserted specific applicability of portions off the published application of Moore et al. against specific ones of applicant's claims. In addition to the fact that the rejection of applicants claim 1 based upon this cited document cannot stand for the reasons cited above, it is also clear that with respect to other ones of applicants claims, the rejection is also unsustainable. Accordingly, specific ones of applicants claims are now addressed below.

Claim 3

Claim 3 depends on the definition of "session node." A session is a specific DMAPI concept defined in the standard. Prior patents, assigned to the same assignee as the present invention, in this area describe this in detail. See explicitly the following quote from US Patent 6,990 478: "The node on which the session is created is designated as the session node, and all specified events generated by file system operations are reported to the session node, regardless of the node at which the events are generated." The Examiner's concept of a "session" is taken from the network world rather than the data management world.

Claim 4

Claim 4 asserts the ability to limit actions which are permitted by the data management rights to specific nodes for the purpose of transferring data. The cited section of Moore [0047]-[0048] does not address data management rights in any way. It addresses user application rights to the data and this is within the context of an implementation that does implement DMAPI. See Figure 7 of Moore. Step 96 acquires a token for the sole purpose of insuring that it knows whether to initiate event processing based on the latest state of the system. All the event processing is done on the metadata server in their step 138 and the only node with data management rights is that metadata server. It does not delegate those rights to any other node which is the most significant feature of the present invention.

Claim 9, 10

Claim 9 (along with claims 7-12 as well) include recitations of the basic functions of the prior patents on which the present application is built. The Examiner rejects claim 9 by citing the same lines of Moore et al. as he did for claim 4 and again, those lines are not at all applicable to data management rights in the DMAPI sense. In fact, [0048] explicitly calls for multiple metadata servers for multiple file systems: Other node(s) may serve as metadata server(s) for other file systems. A metadata server here deals with one file system and there is no multinode capability. The same argument holds for the Examiner's rejection of claim 10.

Claim 11

The Examiner's rejection of claim 11 is again a question of context. Cited portion [0049] refers to the vnode operations which support the POSIX standard for all file systems. The present disclosure operates in the context of a initiating vnode operation also, but that's not the subject of this invention. The DMAPI standard calls for additional forms of moving data into or out of files and punching holes in files which operate within the context of the operations cited in [0049]. If, for example, an application chooses to read a file which is not resident within a file system, the file system must initiate a process which retrieves the data from an external source and writes it into the file system using these additional DMAPI operations. When the DMAPI write is complete, the read may proceed. Portion [0049] does not address these DMAPI calls. Portion [0072]-[0077] of the Moore et al. application appears to describe these operations; however, they address a single file system and a single node (the metadata server for the file system) performing those operations. The present invention has added parallelism, a structure which is not recited or even considered in the application of Moore et al.

Claim 16

Claim 16 describes a key feature of the present invention. There is nothing in the application of Moore et al. which creates worker threads for the processing of DMAPI events or allows those threads to be dispatched on nodes other than the metadata server.

Claim 17, 18

Claim 17 and Claim 18 involve handling of failures. Moore et al. have coupled metadata serving and DMAPI event handling very tightly and provided a mechanism for failover of metadata handling. They have, however, not described in any way what happens to active DMAPI sessions when a metadata server fails. One could conjecture that they do something in this area, but there is nothing to indicate that they have the level of transparency to the user application that is provided by the present invention.

For all of the reasons indicated above, it is seen that the rejection of applicants claims under 35 USC § 102 based upon the published patent application all of Moore at all cannot be sustained. It is therefore very respectfully requested that this rejection be withdrawn.

Sufficiency of the affidavit

The Examiner has also asserted that the previously submitted affidavit under 37 CFR § 1.131 is inadequate. Applicants vehemently disagree with this assertion. It reflects the Examiner's lack of knowledge in the area of software design and development. This is a different area of expertise than one finds in the mechanical arts where a reduction to practice of a golf tee or Frisbee is in issue. In the software arts, code is written and applied against a large plurality of scenarios and conditions to test its workability. **When that workability has been accomplished the result is not a piece of wood or piece of plastic but the successful completion of a process.**

In the present application this process was evidenced by the successful accomplishment of data movement which is the stated objective of the claimed process. What would the Examiner have us provide in order to demonstrate that data was moved from one location to another? Are such applicants forced to thereby submit to the patent office whole disk drives with dates and data thereon?

In any enterprise which commercially produces a software product there is a process for vetting that software. That process can be likened to business records production. Accordingly, the present Applicants assert that the submitted affidavit describes a software vetting process wherein the results of that process was a successful test of that software. In software testing, the

POU920030196US1

proof of the pudding, so to speak, is the passage of that software to a stage of commercialization which by itself is indicative of a successful completion of software test. In short, the present Applicants have set forth a process by which all software is tested, not just by the present assignee, but by every significant commercial software producer. Put another way, the present Applicants have demonstrated that there is a process for maintaining a business record pertaining to software development, testing and completion. Furthermore, Applicants' affidavit clearly sets forth the fact that the present software was tested and completed and that records of this completion were created in accordance with a well-established business records rule.

Accordingly, it is seen that, in the context of software development, the presently submitted affidavit is ample evidence of a successful completion of the invention as of the date indicated. It is therefore respectfully requested that the affidavit be considered to be fully indicative of a successful proof of concept and a successful reduction to practice. Considered as such, it is seen that the reduction of practice of the present invention prior to the priority date of the patent application of Moore et al. is more than sufficient to remove it as a reference in the present prosecution. For this reason also it is seen that the rejection of applicants claims based upon the cited published patent application of Moore et al. be withdrawn.

For all the reasons asserted above, it is Applicants' position that the rejection of claims 1-4 and 9-20 under 35 USC § 102(e) based upon the published patent application of Moore et al., cannot be sustained. It is therefore respectfully requested that it be withdrawn.

35 USC § 103 Rejection

Attention is now directed to the other rejection imposed by the Examiner. In particular, claims 5-8 also stand as being rejected under 35 USC § 103 based upon the aforementioned Moore et al. application in further review of the published patent application of Dugan et al.

The Dugan application is a networking application which the Examiner appears to be citing because all of their use of the word "session." However, it is noted that this term has dual meanings in the computer arts. These disparate meanings are described above but it is noted that

in the present context, as used by a Dugan, the word “session” has nothing at all to do with data management which is the subject of the present invention.

With specific reference to claim 5, it is noted that the word “session identifier” is to be found therein. It is noted that Moore et al. also appeared to teach the use of a session identifier. However, they do so only because the standard requires one. Apart from that it has no use in the cited document. The added patentable value set forth in claim 5 is that the session identifier has meaning on multiple threads on multiple nodes. **Nothing in Moore et al. goes in that direction and Dugan has a different concept of session.**

With specific reference to claim 6, it is seen that it employs the use of a session identifier across multiple nodes. **Again, nothing in Moore et al. allows a DMAPI session to span nodes and nothing in Dugan refers to DMAPI sessions.**

It is also pointed out with respect to claims 7 and 8, that they recite the use of **DMAPI sessions**. In contrast, the Examiner's rejection refers to **network sessions**.

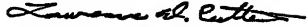
For all these reasons it is seen that the rejection of Applicants' claims based upon the combination of the two cited documents cannot be sustained. Again, it is therefore very respectfully requested that the rejection of Applicants' claims under 35 USC § 103 be withdrawn as well.

It is noted that the present response does not require the payment of any fee. From the above, it is seen that the rejections of Applicants' claims 1-20 cannot be sustained. It is therefore respectfully requested that they be withdrawn. Is also noted that the present response is being submitted within two months of the date of the Final Rejection. Accordingly, it is therefore respectfully requested that the Examiner provide Applicants with an advisory action prior to May 20, 2008.

Applicants' attorney wishes to point out to the Examiner that should the Examiner find it either necessary or desirable to discuss the present response with the claims, that he would be willing to discuss any matter which would assist the Examiner in furthering the prosecution of the present application with all due and proper regard for the appropriate scope of the invention.

Accordingly, should the Examiner wish to discuss this case with Applicants' attorney, the Examiner may contact Applicants' attorney via any of the numbers listed below.

Respectfully submitted,



Lawrence D. Cutter
Attorney for Applicants
Registration No.: 28,501

Dated: April 18, 2008

HESLIN ROTHENBERG FARLEY & MESITI P.C.
5 Columbia Circle
Albany, New York 12203-5160
Telephone: (518) 452-5600
Direct Number: (845) 255-3276
Facsimile: (518) 452-5579